

Learning to Hire Teams*

Adish Singla

ETH Zurich
adish.singla@inf.ethz.ch

Eric Horvitz

Microsoft Research
horvitz@microsoft.com

Pushmeet Kohli

Microsoft Research
pkohli@microsoft.com

Andreas Krause

ETH Zurich
krausea@ethz.ch

Abstract

Crowdsourcing and human computation has been employed in increasingly sophisticated projects that require the solution of a heterogeneous set of tasks. We explore the challenge of building or hiring an effective team, for performing tasks required for such projects on an ongoing basis, from an available pool of applicants or workers who have bid for the tasks. The recruiter needs to learn workers' skills and expertise by performing online tests and interviews, and would like to minimize the amount of budget or time spent in this process before committing to hiring the team. How can one optimally spend budget to learn the expertise of workers as part of recruiting a team? How can one exploit the similarities among tasks as well as underlying social ties or commonalities among the workers for faster learning? We tackle these decision-theoretic challenges by casting them as an instance of online learning for best action selection. We present algorithms with PAC bounds on the required budget to hire a near-optimal team with high confidence. Furthermore, we consider an embedding of the tasks and workers in an underlying graph that may arise from task similarities or social ties, and that can provide additional side-observations for faster learning. We then quantify the improvement in the bounds that we can achieve depending on the characteristic properties of this graph structure. We evaluate our methodology on simulated problem instances as well as on real-world crowdsourcing data collected from the oDesk platform. Our methodology and results present an interesting direction of research to tackle the challenges faced by a recruiter for contract-based crowdsourcing.

Introduction

The success of a project or a collaborative venture depends critically on acquiring a team of contributors. Beyond increased performance and productivity, hiring a strong team leads to enhanced engagement and retention of workers.

"A small team of A+ players can run circles around a giant team of B and C players." – Steve Jobs

Crowdsourcing and outsourcing via online marketplaces further underscores the promise of developing procedures for identifying potential contributors and compos-

ing teams. Crowdsourcing and human computation platforms highlight the opportunities for optimizing team building even when a job requester and workers may be half a world apart and have no advance contact. To date, online crowdsourcing markets have largely focused on micro-tasking through enlisting a non-expert crowds of workers, who work independently and contribute to the solution of simple tasks such as performing image annotation and rating web pages. With the increasing complexity of tasks that are crowdsourced, as well as enterprises outsourcing their work, the need to hire skilled workers with an eye to considerations of complementarity and coordinative efforts in a collaboration around problem solving is becoming important. Contract-based crowdsourcing is another emerging paradigm where workers are recruited on a contract for performing tasks on an ongoing basis. The online platforms are offering new capabilities to deal with this shift towards expertise-driven crowdsourcing. For instance, oDesk provides opportunities for workers to do self-assessments via the taking of voluntary tests ranging from those evaluating language skills to competencies in more complex disciplines such as programming. The platform provides support for recruiters to conduct interviews and perform online tests for job applicants. Furthermore, most of these marketplaces employ a feedback mechanism that allows task and platform owners to track the skill-specific expertise and reputation of workers to help with future recruiting.

Tasks and the team. We consider the crowdsourcing setting where the job requester has a predefined heterogeneous set of types of tasks that need to be solved on an ongoing basis. The notion of task types here could alternatively be taken to refer to the unique set of skills that are needed for addressing the needs of a project. For instance, consider an enterprise whose goal is to outsource a project that has three components or categories of tasks, each requiring a particular skill: (i) web development, (ii) English to Spanish translation, and (iii) video editing. The project would have ongoing assignments of tasks that would belong to one of these three components. When a new task needs to be executed, it is assigned to the hired team and can be performed by the worker possessing the highest expertise for the skill required for this task. The quality of the hired team could then be quantified by the highest expertise among the team members for each of the skills that are required for this project.

*Short version of this paper will appear in HCOMP'15.
Copyright © 2015, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Illustration of approach on toy example with five workers and three (types of) tasks.

Learning workers’ expertise. In the general case, workers’ expertise over different types of tasks or the skills is unknown to the recruiter. To learn the worker’s expertise for a given type of task, the recruiter can perform an online test or evaluate the performance of the worker via assignment of gold-standard questions for which the ground truth is available. Under standard statistical assumptions, performing more of these tests on a worker would give a better estimate about the expertise level of the worker for a given type of task. The recruiter’s goal is to hire a near-optimal team with high likelihood. The main research question is then how to optimally spend the budget (or minimize the total number of tests performed) in order to obtain a sufficiently good estimate of the workers’ expertise over all of the required task types and to be able to make the hiring decision under an allowed level of error tolerance.

Exploiting commonalities. Typically, the unique number of task types and the total number of job applicants (or the workers that bid for the posted tasks) could be large and hence may require performing large numbers of tests in order to learn the workers’ expertise. However, in order to speed up learning, one may be able to exploit the similarities among the tasks and underlying social ties or commonalities among the workers. For instance, consider two types of tasks, requiring skill “java script” and skill “ajax”. By using group testing, the recruiter may design one test for skill “java script” that could allow to additionally infer the expertise on skill “ajax” at no additional cost. Prior knowledge about correlations among workers’ expertise and workers’ features (such as demographics) could also be exploited. Depending on the specific application setting, one may be able to exploit the social ties among workers (or “participants”). The goal is to design algorithms that can exploit these different kinds of commonalities should they be present.

Our contributions can be summarized as follows:

- We present an algorithmic approach to hiring a team of workers as faced by a recruiter for contract-based crowdsourcing;
- we provide algorithms with PAC bounds on the required budget to hire a near-optimal team with high confidence. Our algorithms phrase the decision-theoretic problem of team hiring as an instance of online learning for best action selection.
- We propose a simple model to jointly consider the commonalities among tasks and workers, extend our algorithms to exploit them and
- evaluate the proposed methods using synthetic data as well as data collected from the oDesk platform.

Related Work

Heterogeneous crowdsourcing markets. Our work tackles challenges that arise in heterogeneous crowdsourcing markets where a worker’s performance for a given task depends on the required skills and the expertise level of the worker for those skills. Difallah, Demartini, and Cudré-Mauroux (2013) focus on building automated tools to pick the right set of eligible workers for a given task based on the social networking profile of the workers. Goel, Nikzad, and Singla (2014) design a mechanism for assigning tasks to workers, under the constraints given in terms of a bipartite graph capturing skills and expertise compatibility of the tasks and workers. Another line of research in these markets involve the study of coordination among workers and formation of teams to perform a desired task. Shahaf and Horvitz (2010) introduce the notion of generalized task markets, and how machines and humans can interact together to solve such generalized tasks by forming teams. Zhang et al. (2011) discuss human computation tasks that require effective coordination among workers, such as itinerary planning or data sorting. Our work presents an algorithmic approach to the challenge of team hiring, with guarantees on optimality of the team and the budget required.

Learning in crowdsourcing. Many problems about learning the performance and characteristics of the crowd can be cast as an instance of online learning with associated explore-exploit dilemma, and hence several solutions use the framework of multi-armed bandits (MAB) (Lai and Robbins 1985). Ho and Vaughan (2012) and Ho, Jabbari, and Vaughan (2013) tackle the algorithmic questions concerning learning worker’s expertise, task assignment and label inference for heterogeneous classification tasks. However, their goal is to improve the overall prediction accuracy at lower cost, rather different from our work. Singla and Krause (2013) and Ho, Slivkins, and Vaughan (2014) consider budgeted variants of MAB for learning the price curve and dynamically adjusting payments based on the quality.

Best action selection. From a technical perspective, the most similar work to ours is the best action selection problem, a more recently introduced variant of MAB problems (Even-Dar, Mannor, and Mansour 2006; Bubeck, Munos, and Stoltz 2009; Kalyanakrishnan et al. 2012; Chen et al. 2014). In these settings, the principal agent explores the problem space (the set of actions or “arms”) for a certain time or budget, commits to a policy of the actions and then exploits. Even-Dar, Mannor, and Mansour (2006) study this model under the PAC (probably approximately correct) setting (Valiant 1984) and introduce various (ϵ, δ) -PAC algo-

gorithms for best “arm” identification, *i.e.*, provide bounds on the number of samples required to output an ϵ -optimal action with probability at least $(1 - \delta)$ using concentration bounds (Hoeffding 1963). Kalyanakrishnan et al. (2012) design an adaptive (ϵ, δ) -PAC algorithm LUCB-1 for selecting m best actions, using upper and lower confidence bounds. Zhou, Chen, and Li (2014) also study the problem of selecting m best actions, introducing a new aggregate metric and then applying it to the crowdsourcing setting by simulation experiments. The uniform exploration policy introduced by Even-Dar, Mannor, and Mansour (2006) and the adaptive policy LUCB-1 of Kalyanakrishnan et al. (2012) are the main building blocks of our proposed algorithms. Gabillon et al.; Wang, Viswanathan, and Bubeck (2011; 2013) consider the problem of best arm identification in multiple MAB problem instances by jointly learning over all the problem instances. Our algorithms are also inspired from this idea of jointly identifying best actions over multiple problem instances and we extend the LUCB-1 algorithm to this setting.

Exploiting commonalities and modeling side-observations. A recent line of research has introduced the notion of side-observations to exploit the additional information that can speed up learning. Mannor and Shamir (2011) consider a class of problems that interpolate between bandit feedback and full information settings. They consider the bandit feedback model with side-observations (for instance, such side-observations could arise from user/advertisement similarity, sensor proximity *etc.*) and design algorithms for adversarial settings. Caron et al. (2012) and Buccapatnam, Eryilmaz, and Shroff (2014) extend the results of the side-observation model for stochastic settings. Side-observations through correlations (Fang and Tao 2014) captures bandit problems where the actions are correlated, as well as pulling one actions invokes these correlated actions accounting for additional rewards and observations, motivated by applications in social advertisement. Cesa-Bianchi, Gentile, and Zappella (2013) present an algorithm for contextual bandits correlated through an underlying graph. We borrow some of the ideas from Mannor and Shamir (2011) and Buccapatnam, Eryilmaz, and Shroff (2014) to exploit the commonalities among tasks and among workers. We present a simple model to jointly consider the commonalities among tasks and workers by representing it as a cross product of two side-observation graphs. Furthermore, for the first time, we apply these side-observation models to the best action selection problem.

Problem Statement

We now formalize the problem addressed in this paper.

Tasks and workers. We have a set of M types of tasks (simply referred to as tasks henceforth) and N workers (or job applicants) denoted by the sets $\mathcal{O} = \{o_1, o_2, \dots, o_M\}$ and $\mathcal{W} = \{w_1, w_2, \dots, w_N\}$, respectively. We shall assume $N \geq M$, simply meaning that there is at least one unique job applicant per type of task. For instance, in Figure 1, we have $M = 3$ tasks, and $N = 5$ workers. We model the performance of a worker for a given task as a bounded random variable with unknown mean. Assigning task $o_j \in \mathcal{O}$ to worker $w_i \in \mathcal{W}$ at time t yields a performance value (as

feedback) denoted by random variable $X_{(i,j)}^t$, sampled from an unknown distribution with mean value $\mu_{(i,j)}$. For simplicity and w.l.o.g, we shall assume that the underlying distribution from which $X_{(i,j)}^t$ is sampled has a bounded support within $[0, 1]$. The mean performance values are denoted by an unknown performance matrix $\mu : N \times M \rightarrow \mathcal{R}_{\geq 0}$ with tasks as columns, and workers as rows. We assume a stochastic setting where $X_{(i,j)}^t$ are *i.i.d.* for any fixed pair of worker w_i and task o_j . Also, $X_{(i,j)}^t$ are *independent* across i, j and t .

Side-observation model. The workers and tasks are embedded in some (known) underlying graphs, denoted by $G_w(V_w, E_w)$ and $G_o(V_o, E_o)$. The nodes $V_w \in G_w$ correspond to the N workers, and nodes $V_o \in G_o$ correspond to the M tasks. We shall assume undirected graphs, though the models and results could be extended to the setting of directed graphs as well. The edges in these graphs capture the model of side-observations that may be possible to obtain at no additional cost (Mannor and Shamir 2011; Buccapatnam, Eryilmaz, and Shroff 2014). In our model, when worker w_i is assigned task o_j at time t , apart from observing the performance $X_{(i,j)}^t$, the following additional set of observations become available:

- $X_{(i,q)}^t \forall q : \{o_j, o_q\} \in E_o$, the additional observations associated with the tasks neighboring to o_j in G_o .
- $X_{(p,j)}^t \forall p : \{w_i, w_p\} \in E_w$, the additional observations associated with the workers neighboring to w_i in G_w .

In Figure 1, assigning task o_1 to worker w_2 at time t would yield set of observations given by $X^t = \{X_{(2,1)}^t, X_{(2,2)}^t, X_{(1,1)}^t, X_{(3,1)}^t\}$. The goal is to design algorithms that can exploit these side-observations *whenever* present, and smoothly interpolate between the *bandit setting* (absence of side-observations, $E_w = E_o = \emptyset$) to the *full information setting* (fully connected graphs).

The objective. Our goal is to select or hire a team of workers denoted by \mathcal{S}^* , of size at most M from the set \mathcal{W} , comprising the highest performing worker for each task $o \in \mathcal{O}$. If the performance matrix is known, the problem is trivial, for instance, in Figure 1, the optimal team is $\{w_1, w_2, w_4\}$. Hence, the goal is to design algorithm that can efficiently learn the performance matrix $\mu[N, M]$ and output a near-optimal team. In our model, a team \mathcal{S} is ϵ -optimal, when, for each task $o_j \in \mathcal{O}$, we have:

$$\forall o_j \in \mathcal{O}, \max_{w_i \in \mathcal{W}} \mu_{(i,j)} - \max_{w_i \in \mathcal{S}} \mu_{(i,j)} \leq \epsilon \quad (1)$$

In Figure 1, $\{w_1, w_3, w_4\}$ is an ϵ -optimal team for $\epsilon = 0.01$. Given our stochastic assumptions, the algorithm can repeatedly assign a task o_j to worker w_i in order to get a good estimate of the performance $\mu_{(i,j)}$. We call each assignment being a *test* performed. We assume that each such test poses a unit cost to the algorithm. We seek algorithms with PAC bounds, *i.e.*, for given positive constants (ϵ, δ) , the algorithm should output an ϵ -optimal team with probability of at least $(1 - \delta)$. We measure the efficiency of such an algorithm in terms of the total number of tests required or equivalently the budget spent.

Algorithms for Budgeted Hiring

Overview of basic approach

To present some of the key insights in designing our algorithms, we first consider a simple setting.

Single task ($M = 1$) without side-observations. Let us first consider the simple setting of hiring to solve one task, *i.e.*, $M = 1$ and the goal is to find an ϵ -optimal worker from set \mathcal{W} with success probability of at least $(1 - \delta)$. We consider the recruiting of team members from among N workers as the set of actions at hand, and reduce the decision problem to the problem of best action selection (Even-Dar, Mannor, and Mansour 2006; Bubeck, Munos, and Stoltz 2009; Kalyanakrishnan et al. 2012; Chen et al. 2014). For example, the NAIVE(ϵ, δ) algorithm of Even-Dar, Mannor, and Mansour (2006) provides (ϵ, δ) -PAC guarantees by uniformly allocating a sufficient number of observations for each action to be able to select ϵ -optimal action with probability at least $(1 - \delta)$. By using Hoeffding’s inequality (Hoeffding 1963), a sufficient number of observations is $\lceil \frac{2}{\epsilon^2} \ln(\frac{N}{\delta}) \rceil$. This NAIVE(ϵ, δ) algorithm is the main building block for our proposed algorithm UEXPSELECT based on uniform exploration of the actions.

This algorithm is based on uniform exploration and ignores the fact that some actions may be easier or harder to distinguish. For example, in Figure 1, considering task o_1 , distinguishing w_5 from w_1 is easier than distinguishing w_2 from w_1 . To tackle this problem, Kalyanakrishnan et al. (2012) design an adaptive (ϵ, δ) -PAC algorithm LUCB-1 using upper and lower confidence bounds. LUCB-1 adapts to the complexity of the problem instance, provides distribution dependent bounds and is the state-of-the-art algorithm for the best action selection problem. We use LUCB-1 as the main building block for our proposed algorithm AEXPSELECT, an adaptive variant of UEXPSELECT.

Multiple tasks ($M > 1$) without side-observations. One possible way to tackle this challenge is to consider each task as a separate instance of the best action selection problem, and to use one of the previously discussed algorithms NAIVE or LUCB-1 separately. However, one can hope to do better by jointly considering all of the tasks, and allocating the budget across tasks in an adaptive manner. For instance, in Figure 1, the task t_2 is harder than task t_1 and task t_3 in terms of distinguishing and selecting the best worker. Recently, Gabillon et al. (2011) and Wang, Viswanathan, and Bubeck (2013) have addressed this problem of best arm identification in multiple multi-armed bandit (MAB) instances by jointly learning over all of the instances. Our proposed algorithms are inspired from the idea of jointly identifying best workers (the team) for all of the tasks and AEXPSELECT extend the LUCB-1 algorithm to this setting.

Exploiting side-observation graphs

Side-observation models (Mannor and Shamir 2011; Caron et al. 2012; Buccapatnam, Eryilmaz, and Shroff 2014) have been studied mainly in context of regret minimization problems using MAB framework, modeling the observations via an underlying graph connecting the “arms” of the MAB. Although different ideas have been explored on how to ex-

ploit side-observations via an underlying graph, all these ideas revolve around the *minimal dominating set* of the side-observation graph $G(V, E)$ denoted by $\text{DOM}(G)$. This concept refers to the smallest subset of vertices that cover the rest — every vertex of the graph G is either in $\text{DOM}(G)$ or is directly connected to one of the vertices in $\text{DOM}(G)$.

We extend these ideas to apply the side-observation models for the best action selection problem. Since our proposed algorithms jointly learn over these M tasks, we would like to jointly exploit the side-observation graphs over the tasks and the workers. We can model the side-observation graphs jointly as the cartesian product of two graphs given by $G_w \square G_o$, denoted as $G_{wo} = (V_{wo}, E_{wo})$. In a cartesian product of graphs, the vertices are given by the cartesian product of the vertex sets of the individual graphs, $V_{wo} = V_w \times V_o$, or alternatively, $V_{wo} = \{(w_i, o_j) : i \in [1 \dots N] \text{ and } j \in [1 \dots M]\}$, *i.e.*, G_{wo} has $M \cdot N$ vertices. The edges are given by E_{wo} such that $\{(w_i, o_j)\}$ and $\{(w_{i'}, o_{j'})\}$ have an edge if either *i*) $w_i = w_{i'}$ and $(o_j, o_{j'}) \in E_o$, or *ii*) $o_j = o_{j'}$ and $(w_i, w_{i'}) \in E_w$. Let $\gamma(G_{wo})$ denote the minimum size of a dominating set in the resulting graph. Computing the dominating set itself is NP-Hard by a reduction from the set-cover problem (Guha and Khuller 1998). However, an approximate solution can be found of size upper bounded by $(1 + \ln(1 + \text{DEG}(G_{wo})))$, where DEG denotes the maximum degree of any vertex in the graph (Guha and Khuller 1998). Let us denote this approximate dominating set as $\widetilde{\text{DOM}}(G_{wo})$ and the corresponding approximate dominating number as $\tilde{\gamma}_{G_{wo}}$.

We denote the set of *actions* as $\mathcal{A} = \{a_{(i,j)} : i \in [1 \dots N] \text{ and } j \in [1 \dots M]\}$. Taking action $a_{(i,j)}$ at time t is equivalent to assigning a worker w_i to a task o_j at time t . For any action $a_{(i,j)}$, with a slight abuse of notation, we denote its neighboring action belonging to the dominating set as $\widetilde{\text{DOM}}(G_{wo}, a_{(i,j)}, \cdot) \in \widetilde{\text{DOM}}(G_{wo})$. We call this the *dominating action* for $a_{(i,j)}$. For any action $a_{(i,j)} \in \widetilde{\text{DOM}}(G_{wo})$, we denote the set of actions dominated by $a_{(i,j)}$ as $\widetilde{\text{DOM}}(G_{wo}, \cdot, a_{(i,j)}) \subseteq V_{wo}$. The main idea used in our algorithms UEXPSELECT and AEXPSELECT is to replace the picked action $a_{(p,q)}^t$ by its dominating action $a_{(\bar{p}, \bar{q})}^t$.

Model parameters and execution variables

We now introduce several model parameters as well as notation that will be useful to describe the algorithms.

ϵ -optimal team. For any task o_j , the highest performance among all the workers is given by $\mu_{(i^*, j)} = \max_{w_i \in \mathcal{W}} \mu_{(i, j)}$ and let $w_{(i^*, j)}$ be the worker with this highest performance. We denote the best worker for this task with the corresponding action $a_{(i^*, j)} \in \mathcal{A}$. For a given task o_j , we can now denote the relative quality of a given worker w_i w.r.t the performance of the best worker available for this task as $\Delta_{(i, j)} = \mu_{(i^*, j)} - \mu_{(i, j)}$. For the specific case of best worker $w_{(i^*, j)}$, this quantity is defined as: $\Delta_{(i^*, j)} = \mu_{(i^*, j)} - \max_{w_i \in \mathcal{W} \setminus \{w_{(i^*, j)}\}} \mu_{(i, j)}$ denoting the gap with the second best worker for this task. For any task o_j , we say a worker $w_i \in \mathcal{W} \setminus \{w_{(i^*, j)}\}$ is ϵ -optimal for o_j if $\Delta_{(i, j)} \leq \epsilon$. We denote this set of ϵ -optimal workers along with best worker $w_{(i^*, j)}$ as $S_{(\epsilon, j)}$. Now, a solution out-

Algorithm 1: Algorithm UEXPSELECT

1 **Input:** Tasks: \mathcal{O} ; Workers: \mathcal{W} ; Side observation graphs: G_w, G_o ; PAC parameters: (ϵ, δ) ;
2 **Output:** Team of workers $S \subseteq W : |S| \leq M$, such that S is ϵ -optimal with probability at least $(1 - \delta)$;
3 **Initialize:**

- *Compute:* $G_{wo} = G_w \square G_o$; $\widetilde{\text{DOM}}(G_{wo})$;
- $t = 0$; $S = \emptyset$;
- $\forall a_{(i,j)} \in A: \mu_{(i,j)}^t = 0$; $n_{(i,j)}^t = 0$; $y_{(i,j)}^t = 0$;

while $\exists a_{(i,j)} \in A : y_{(i,j)}^t < \lceil \frac{2}{\epsilon^2} \ln(\frac{M \cdot N}{\delta}) \rceil$ **do**

4 $a_{(p,q)}^t = \arg \min_{a_{(i,j)} \in A} y_{(i,j)}^t$; \triangleright Greedy action
5 $a_{(\bar{p},\bar{q})}^t = \widetilde{\text{DOM}}(G_{wo}, a_{(p,q)}^t, \cdot)$;
6 **Perform action:** $a_{(\bar{p},\bar{q})}^t$; \triangleright Assign $o_{\bar{q}}$ to $w_{\bar{p}}$
7 **Feedback:** Obtain observations X^t for actions dominated by $a_{(\bar{p},\bar{q})}^t$; $\widetilde{\text{DOM}}(G_{wo}, \cdot, a_{(\bar{p},\bar{q})}^t)$;
8 **Update Variables:**

- $n_{(\bar{p},\bar{q})}^{t+1} = n_{(\bar{p},\bar{q})}^t + 1$;
- $\forall a_{(i,j)} \in \widetilde{\text{DOM}}(G_{wo}, \cdot, a_{(\bar{p},\bar{q})}^t), y_{(i,j)}^{t+1} = y_{(i,j)}^t + 1$;
- $\forall a_{(i,j)} \in \widetilde{\text{DOM}}(G_{wo}, \cdot, a_{(\bar{p},\bar{q})}^t), \text{update } \mu_{(i,j)}^{t+1} \text{ from } X^t$;
- $t = t + 1$;

foreach $j \in [1 \dots M]$ **do**

9 $w_{(i^*,j)}^t = \arg \max_{w_i \in \mathcal{W}} \mu_{(i,j)}^t$;
10 $S \leftarrow S \cup \{w_{(i^*,j)}^t\}$;

11 **Output:** S

put S from the algorithm is ϵ -optimal (denoted as S_ϵ) if it contains at least one ϵ -optimal worker for each task, i.e., $\forall j \in [1 \dots M] : |S \cap S_{(\epsilon,j)}| \geq 1$. Putting $\epsilon = 0$ in S_ϵ will correspond to the optimal team.

Variables over execution. The algorithm will run in time steps, denoted by t , where each time step corresponds to the assignment of a task to a worker. Hence, the total number of time steps until execution of the algorithm corresponds to the budget spent or sample complexity for the algorithm. At time step t , let $n_{(i,j)}^t$ correspond to the number of times task o_j has been assigned to w_i (or simply, action $a_{(i,j)}$ has been performed). Also, let $y_{(i,j)}^t$ correspond to the number of times total observations have been made about performance of w_i for task o_j (note that, in the absence of side-observations, $n_{(i,j)}^t = y_{(i,j)}^t$). The current estimate of the mean values are denoted by $\mu_{(i,j)}^t$. With these estimates, we also define $\mu_{(i^*,j)}^t = \max_{w_i \in \mathcal{W}} \mu_{(i,j)}^t$, and $w_{(i^*,j)}^t = \arg \max_{w_i \in \mathcal{W}} \mu_{(i,j)}^t$. Similarly, we define the quantities $\Delta_{(i,j)}^t$ based on current estimate of the performance values $\mu_{(i,j)}^t$.

Algorithm UEXPSELECT

We now present our first algorithm UEXPSELECT, shown in Algorithm 1, based on the uniform exploration of all the actions extending ideas of NAIVE algorithm (Even-Dar, Mannor, and Mansour 2006). At each iteration, the algorithm selects the action $a_{(p,q)}^t$ with minimal number of observations

Algorithm 2: Algorithm AEXPSELECT

1 **Input:** Tasks: \mathcal{O} ; Workers: \mathcal{W} ; Side observation graphs: G_w, G_o ; PAC parameters: (ϵ, δ) ;
2 **Output:** Team of workers $S \subseteq W : |S| \leq M$, such that S is ϵ -optimal with probability at least $(1 - \delta)$;
3 **Initialize:**

- *Compute:* $G_{wo} = G_w \square G_o$; $\widetilde{\text{DOM}}(G_{wo})$;
- $t = 0$; $S^t = \emptyset$; $R^t = \mathcal{O}$;
- $\forall a_{(i,j)} \in A: \mu_{(i,j)}^t = 0$; $n_{(i,j)}^t = 0$; $y_{(i,j)}^t = 0$;
- $\forall a_{(i,j)} \in A: \beta_{(i,j)}^t \rightarrow \infty$; $\forall o_j \in \mathcal{O}: \Delta_j^t \rightarrow \infty$;

while $R^t \neq \emptyset$ **do**

4 $o_q^t = \arg \max_{o_j \in R^t} \Delta_j^t$;
5 $w_{(i^*,q)}^t = \arg \max_{w_i \in \mathcal{W}} \mu_{(i,q)}^t$;
6 $w_{(i^*,q)}^t =$
 $\arg \max_{w_i \in \mathcal{W} \setminus \{w_{(i^*,q)}^t\}} (\mu_{(i,q)}^t + \beta(y_{(i,q)}^t, t))$;
7 $w_p^t = \arg \max_{w_i \in \{w_{(i^*,q)}^t, w_{(i^*,q)}^t\}} \beta(y_{(i,q)}^t, t)$;
8 $a_{(p,q)}^t \leftarrow (w_p^t, o_q^t)$; \triangleright Greedy action
9 $a_{(\bar{p},\bar{q})}^t = \widetilde{\text{DOM}}(G_{wo}, a_{(p,q)}^t, \cdot)$;
10 **Perform action:** $a_{(\bar{p},\bar{q})}^t$; \triangleright Assign $o_{\bar{q}}$ to $w_{\bar{p}}$
11 **Feedback:** Obtain observations X^t for actions dominated by $a_{(\bar{p},\bar{q})}^t$; $\widetilde{\text{DOM}}(G_{wo}, \cdot, a_{(\bar{p},\bar{q})}^t)$;
12 **Update Variables:**

- $n_{(\bar{p},\bar{q})}^{t+1} = n_{(\bar{p},\bar{q})}^t + 1$;
- $\forall a_{(i,j)} \in \widetilde{\text{DOM}}(G_{wo}, \cdot, a_{(\bar{p},\bar{q})}^t), y_{(i,j)}^{t+1} = y_{(i,j)}^t + 1$;
- $\forall a_{(i,j)} \in \widetilde{\text{DOM}}(G_{wo}, \cdot, a_{(\bar{p},\bar{q})}^t), \text{update } \mu_{(i,j)}^{t+1} \text{ from } X^t$;
- $t = t + 1$;
- $\forall o_j \in R^t, \text{update } \Delta_j^t$;

13 **Update Solution:**

foreach $o_j \in R^t$ **do**

if $\Delta_j^t \leq \epsilon$ **then**

14 $w_{(i^*,j)}^t = \arg \max_{w_i \in \mathcal{W}} \mu_{(i,j)}^t$;

15 $S^t = S^t \cup \{w_{(i^*,j)}^t\}$;

16 $R^t = R^t \setminus \{o_j\}$;

17 **Output:** S^t

$y_{(p,q)}^t$ (Step 4). This choice is natural and can be thought of as “greedy” in order to quickly move towards termination of the algorithm. Given the side-observation model, algorithm takes the action $a_{(\bar{p},\bar{q})}^t$ (Step 5), i.e., the one that dominates $a_{(p,q)}^t$, as taking $a_{(\bar{p},\bar{q})}^t$ also gives us the desired observation needed for $a_{(p,q)}^t$. Then, it receives the observation set X^t , corresponding to all the actions that are dominated by $a_{(\bar{p},\bar{q})}^t$, and updates the corresponding variables. Once every action has made observations of at least $\lceil \frac{2}{\epsilon^2} \ln(\frac{M \cdot N}{\delta}) \rceil$, the algorithm selects the best set S based on the observed performances $\mu_{(i,j)}^t$. Note that, if we ignore the side-observation model, then $a_{(\bar{p},\bar{q})}^t \equiv a_{(p,q)}^t$, and the observations set correspond to singleton set, given by $X^t = \{X_{(p,q)}^t\}$.

Algorithm AEXPSELECT

In order to adapt the algorithm to the variability of the hardness of the problem in identifying suboptimal workers across tasks and within one given task, we present a second algorithm AEXPSELECT, based on ideas of LUCB-1 algorithm (Kalyanakrishnan et al. 2012). In order to present AEXPSELECT, we introduce some specific terminology as well as the approach used to pick the actions.

First, we associate confidence bounds, *i.e.*, a high probability bound over the estimates of the performance $\mu_{(i,j)}^t$. This is denoted by the function $\beta(y_{(i,j)}^t, t)$. The specific form of function we use, as used in LUCB-1, is given by

$\beta(y, t) = \sqrt{\frac{1}{2 \cdot y} \ln \left(\frac{5}{4} \cdot \frac{M \cdot N}{\delta} \cdot t^4 \right)}$. One of the key intuitions behind this specific function is that we seek to ensure that the probability of the event that the confidence interval bounds are *ever* violated over the *lifespan* of the algorithm is bounded by δ . For a given action $a_{(i,j)}$, the upper and lower confidences over the performance estimate $\mu_{(i,j)}^t$ are given as $(\mu_{(i,j)}^t + \beta(y_{(i,j)}^t, t))$ and $(\mu_{(i,j)}^t - \beta(y_{(i,j)}^t, t))$ respectively.

At a given time t and for a given task o_j , we denote the worker with highest empirically observed performance as $w_{(i^*,j)}^t$, given by:

$$w_{(i^*,j)}^t = \arg \max_{w_i \in \mathcal{W}} \mu_{(i,j)}^t \quad (2)$$

Next, from the remaining $N - 1$ workers, we find the worker with maximum value of upper confidence of performance estimate as follows:

$$w_{(i^\bullet,j)}^t = \arg \max_{w_i \in \mathcal{W} \setminus \{w_{(i^*,j)}^t\}} (\mu_{(i,j)}^t + \beta(y_{(i,j)}^t, t)) \quad (3)$$

The empirical mean of $w_{(i^*,j)}^t$ is denoted by $\mu_{(i^*,j)}^t$, and has lower confidence bound of $(\mu_{(i^*,j)}^t - \beta(y_{(i^*,j)}^t, t))$. For $w_{(i^\bullet,j)}^t$, the empirical mean is denoted by $\mu_{(i^\bullet,j)}^t$, and has upper confidence bound of $(\mu_{(i^\bullet,j)}^t + \beta(y_{(i^\bullet,j)}^t, t))$. The quantity that is of particular interest is the gap between upper confidence bound on $\mu_{(i^\bullet,j)}^t$ and lower confidence bound on $\mu_{(i^*,j)}^t$. Intuitively, as we get increasing numbers of observations and confidence widths shrink, this gap should reduce to below zero. We denote this quantify for task o_j as follows:

$$\Delta_j^t = (\mu_{(i^\bullet,j)}^t + \beta(y_{(i^\bullet,j)}^t, t)) - (\mu_{(i^*,j)}^t - \beta(y_{(i^*,j)}^t, t)) \quad (4)$$

Based on the ideas from LUCB-1, the algorithm can commit to worker $w_{(i^*,j)}^t$ for task o_j whenever $\Delta_j^t \leq \epsilon$, and this is ϵ -optimal choice, as long as the confidence intervals are not violated. Intuitively, we are taking the worst-case estimate of $w_{(i^*,j)}^t$ and highest of the best-case estimate from the remaining workers — ensuring this difference being less than ϵ is sufficient to commit to worker $w_{(i^*,j)}^t$.

AEXPSELECT is shown in Algorithm 2. At each iteration, the algorithm first selects the task with highest Δ_j^t denoted by index o_q^t (Step 4). Then, it finds the corresponding

workers $w_{(i^*,q)}^t$ and $w_{(i^\bullet,q)}^t$ (Step 5, 6). Then, the greedy choice of action $a_{(p,q)}^t$ is based on choosing the worker with higher confidence width among $w_{(i^*,q)}^t$ and $w_{(i^\bullet,q)}^t$ (Step 7, 8). Note that, the solution set S^t is built over time. The algorithm maintains a set of tasks R^t as the tasks for which a worker still needs to be selected. As soon as the condition in Equation 4 is met for a task, that task is no longer considered for further actions and removed from R^t . The algorithm terminates when R^t is empty. Note that there is a common time clock across all the tasks. Jointly learning over all the tasks ensures that the algorithm can allocate more assignments for the tasks which have maximum uncertainty. Furthermore, it allows us to jointly exploit the side-observation graphs.

Performance Analysis

We now analyze the performance of the proposed algorithms UEXPSELECT and AEXPSELECT. Most of the results below can be derived using the proof techniques of NAIVE (Even-Dar, Mannor, and Mansour 2006) and LUCB-1 (Kalyanakrishnan et al. 2012), and can be seen as extension of their results.

Performance Bounds for UEXPSELECT

Let us consider the case of the absence of side-observations, which is equivalent to setting $E_w = \emptyset$ and $E_o = \emptyset$. In this case, $a_{(\bar{p},\bar{q})}(t) \equiv a_{(p,q)}^t$, observations set correspond to singleton set of $X^t = \{X_{(p,q)}^t\}$, and $\widetilde{\text{DOM}}(G_{wo})$ and $\text{DOM}(G_{wo})$ are both equal to V_{wo} . In fact, in terms of performance bounds, the algorithm UEXPSELECT can be seen as equivalent to running M instances of NAIVE- $(\epsilon, \frac{\delta}{M})$. Based on Theorem 6 from Even-Dar, Mannor, and Mansour (2006), the sample complexity of NAIVE- (ϵ, δ) for one instance of the problem with N actions is given by $\left(N \cdot \left\lceil \frac{2}{\epsilon^2} \ln \left(\frac{N}{\delta} \right) \right\rceil \right)$. Hence, the sample complexity of UEXPSELECT in the absence of side-observations is given by $\left(M \cdot N \cdot \left\lceil \frac{2}{\epsilon^2} \ln \left(\frac{M \cdot N}{\delta} \right) \right\rceil \right)$. The PAC- (ϵ, δ) guarantees hold simply from the correctness of NAIVE- (ϵ, δ) . The fact that we ran M instances of NAIVE with $\frac{\delta}{M}$ ensures that the error probability is bounded by δ using the union bound. Next, we can state the improvement in performance obtained by accounting for side-observations in Theorem 1.

Theorem 1. *The algorithm UEXPSELECT is (ϵ, δ) -PAC optimal with sample complexity of $\left(\tilde{\gamma}_{G_{wo}} \cdot \left\lceil \frac{2}{\epsilon^2} \ln \left(\frac{M \cdot N}{\delta} \right) \right\rceil \right)$, where $G_{wo} = G_w \square G_o$ and $\tilde{\gamma}_{G_{wo}} \leq (1 + \ln(1 + \text{DEG}(G_{wo}))) \cdot \gamma_{G_{wo}}$.*

Recall that $\widetilde{\text{DOM}}(G_{wo})$ denotes the polynomial-time approximation of the dominating set for G_{wo} and has size bounded by $(1 + \ln(1 + \text{DEG}(G_{wo})))$ (Guha and Khuller 1998). By taking each action of $\widetilde{\text{DOM}}(G_{wo})$ once, the entire set of actions is covered. Hence by taking $\tilde{\gamma}_{G_{wo}}$ actions, we get observations of the $M \cdot N$ actions resulting in a potential saving of tests by factor of $\frac{\gamma_{G_{wo}}}{\tilde{\gamma}_{G_{wo}}}$. Importantly, the greedy way of selecting the actions in Step 4 of Algorithm 1 ensures that all of the actions in $\widetilde{\text{DOM}}(G_{wo})$ are scanned uniformly.

Performance Bounds for AEXPSELECT

Let us define $\Delta_{(i,j),\alpha} = \max\{\Delta_{(i,j)}, \alpha\}$, for any value of $\alpha \geq 0$. In particular, we are interested in quantities $\Delta_{(i,j),\epsilon/2}$. Let us again begin by considering the case of the absence of side-observations. One way to tackle this problem is then to run M instances of LUCB-1- $(\epsilon, \frac{\delta}{M})$ algorithm, each with its own time clock. Based on Theorem 6 from Kalyanakrishnan et al. (2012), the expected sample complexity of LUCB-1- (ϵ, δ) for one instance of the problem with N actions for a particular task o_j is given by:

$$O\left(\left(\sum_{i \in [N]} \frac{1}{\Delta_{(i,j),\frac{\epsilon}{2}}^2}\right) \ln\left(\frac{1}{\delta} \cdot \sum_{i \in [N]} \frac{1}{\Delta_{(i,j),\frac{\epsilon}{2}}^2}\right)\right) \quad (5)$$

The expected sample complexity of running M instances of LUCB-1- $(\epsilon, \frac{\delta}{M})$ is then given by:

$$O\left(\sum_{j \in [M]} \left(\left(\sum_{i \in [N]} \frac{1}{\Delta_{(i,j),\frac{\epsilon}{2}}^2}\right) \ln\left(\frac{M}{\delta} \cdot \sum_{i \in [N]} \frac{1}{\Delta_{(i,j),\frac{\epsilon}{2}}^2}\right)\right)\right) \quad (6)$$

However, by jointly learning across all the tasks, an algorithm can adaptively allocate assignments across the tasks. AEXPSELECT is based on this idea, originally proposed in Gabillon et al.; Wang, Viswanathan, and Bubeck (2011; 2013) and it extends LUCB-1 algorithm to this joint setting. Intuitively, the main reason this is possible in the best-action selection problems is because the problem complexity is defined in terms of relative “gap” $\Delta_{(i,j)}$ which can be mixed together for all the tasks o_j to create one pool of $M \cdot N$ actions defined by their correspond gaps $\Delta_{(i,j)}$. Then, by using a common time clock over these $M \cdot N$ actions, the main technical results of LUCB-1 extends to this joint setting (Wang, Viswanathan, and Bubeck 2013). The sample complexity of AEXPSELECT is given in Theorem 2, which is based on Theorem 6 from (Kalyanakrishnan et al. 2012).

Theorem 2. *The algorithm AEXPSELECT is (ϵ, δ) -PAC optimal with expected sample complexity given by*

$$O\left(\left(\sum_{j \in [M]} \sum_{i \in [N]} \frac{1}{\Delta_{(i,j),\frac{\epsilon}{2}}^2}\right) \ln\left(\frac{1}{\delta} \cdot \sum_{j \in [M]} \sum_{i \in [N]} \frac{1}{\Delta_{(i,j),\frac{\epsilon}{2}}^2}\right)\right)$$

Note that the above sample complexity bound is similar in structure as given in Equation 5 with total of $M \cdot N$. However, this is different compared to one obtained in Equation 6 by running M instances of LUCB-1- $(\epsilon, \frac{\delta}{M})$ with separate time clock for each task. In fact, when all the tasks are of equal hardness defined by quantity $\sum_{i \in [N]} \frac{1}{\Delta_{(i,j),\frac{\epsilon}{2}}^2}$ for a given o_j , the sample complexity in Theorem 2 and Equation 6 is same.

This bound in Theorem 2 is loose in the sense that it doesn’t explicitly account for the performance gain achieved by the side-observations, even though AEXPSELECT uses the same approach as that used in UEXPSELECT to exploit side-observation graphs. In the worst-case, the static model of side-observations (*i.e.*, a pre-computed and fixed

dominating set) that we used, does not help boost the performance for an adaptive algorithm. Intuitively, and as we observed during empirical evaluations, for the problem instances that are uniformly difficult, we tend to gain more value from side-observations. However in such cases, AEXPSELECT tend to behave more closely as UEXPSELECT. For more skewed tasks and workers in terms of difficulty and performance, the “easier” to identify workers and tasks gets “eliminated” over time, and hence the value of side-observations diminish as well. Hence, for adaptive algorithms like AEXPSELECT, a more effective way of exploiting side-observations would need policies that construct dynamic dominating sets at every time step taking into account the remaining uncertainties over the actions.

Experimental Evaluation

We now report on the results of our experiments.

Experimental Setup and Datasets

We compare the performance of adaptive algorithm AEXPSELECT against the uniform exploration based algorithm of UEXPSELECT. Furthermore, we quantify the effect of side-observations by comparing these two algorithms with their variants without side-observation graphs (setting $E_o = \emptyset$, $E_w = \emptyset$ as input).

Metrics and parameters. The primary metric is the quality of the team output by the algorithm for given budget, measured through *i)* average precision, and *ii)* average performance gap, as defined next. For a given output S , and any task o_j , the precision for task o_j is defined to be 1 if S contains an ϵ -optimal worker for task o_j , *i.e.*, $|S \cap S_{(\epsilon,j)}| \geq 1$, else 0. The performance gap for a task o_j is defined to be $(\mu_{(i^*,j)} - \max_{w_i \in S} \mu_{(i,j)})$. We report the average precision and average performance gap over all the M tasks for the team output by the algorithm for a given budget.

The primary quantity that we vary in the experiments is the total number of tests performed or budget spent by the algorithm. For ease of interpretation, we shall use the unit of the average budget spent per worker/task pair. Also, we shall report results by varying the hardness of the problem instance (Figure 2(c)). For a given task o_j , we used the notion of hardness given by $\Delta_j^{min} = \min_{i \in [1 \dots N]} \Delta_{(i,j)}$. We vary average value of the gap Δ_j^{min} over tasks (*i.e.*, $\frac{1}{M} \cdot \sum_{j \in [1 \dots M]} \Delta_j^{min}$) by creating different datasets and measuring the performance of different algorithms for a fixed budget.

The PAC parameters ϵ and δ are fixed for all of the reported experiments and set to 0.05. The number of tasks is $M = 10$ and total number of workers is $N = 200$. In all of the experiments with varying budget, the average Δ_j^{min} over tasks is fixed to 0.25, with Δ_j^{min} for a task o_j uniformly sampled in the range from $[0.01, 0.5]$. For the experiment in Figure 2(c) where the average Δ_j^{min} is varied, the average budget per worker/task pair is fixed to 20, *i.e.*, equivalent to total budget of $M \cdot N \cdot 20$. The values of the performance matrix μ are scaled to lie in the range $[\mu^{min}, \mu^{max}]$ where




Figure 2: Experimental results on synthetic data with absence of side-observation graphs. In Figure 2(a), 2(b), the budget is varied, and the metrics of average precision and average performance gap are measured, respectively. In Figure 2(c), budget is kept fixed to 20 per worker/task pair, and average Δ_j^{min} is changed from 0.25 to 0.05, making the problem instance difficult.




Figure 3: Experimental results on oDesk data. In all the plots, budget is varied and metric of average precision is measured. In Figure 3(a), there is no side-observation graph and is equivalent to the plot in Figure 2(a). Figure 3(b) and Figure 3(c) shows the comparison of AEXPSELECT and UEXPSELECT with their variants without side-observations ($E_o = \emptyset, E_w = \emptyset$ as input).

$\mu^{min} = 0.1$ and $\mu^{max} = 0.9$. We assume a Bernoulli feedback model, i.e., for $\mu_{(i,j)}$, assigning task o_j to worker w_i yields a feedback value of 1 with probability $\mu_{(i,j)}$ and 0 otherwise. All the results are reported as an average of 10 iterations of the algorithms.

Synthetic data. We created synthetic data for $N = 200$ workers and $M = 10$ tasks as follows. For each task o_j , we sampled Δ_j^{min} uniformly at random from range $[0.01, 0.5]$ (to have average $\Delta_j^{min} = 0.25$). Then, to create the performance vector $\mu_{(:,j)}$ for N workers (corresponding to a column in the performance matrix in Figure 1), we sampled $(N - 1)$ values in the range $[\mu^{min}, \mu^{max} - \Delta_j^{min}]$ and one value (of the best worker) is set to μ^{max} . These N values are then randomly permuted and assigned to the N workers for task o_j . This process is repeated for each of the M tasks independently. For the synthetic experiments, we didn't use side-observations, equivalent to having $E_w = \emptyset$ and $E_o = \emptyset$. For the experiment reported in Figure 2(c), we created 4 more variants of the synthetic data by varying average Δ_j^{min} as $[0.25, 0.20, 0.15, 0.10, 0.05]$.

oDesk data. The primary purpose of using data from oDesk is to be able to obtain real-world distributions of the performance matrix, as well as a realistic way of creating the side-observation graphs. oDesk has over 2.7 million freelancers and 0.5 million job requesters worldwide. We used the publicly available API¹ from oDesk to obtain the data

below. In oDesk platform, each posted task or job is assigned to a predefined taxonomy by the job requester. There are 12 top-level categories of the tasks and about 90 second-level categories. We took $M = 10$ tasks, with 4 tasks in the top-level category *Design & Creative*, 3 tasks in the top-level category *Translation* and 3 tasks in *Data Science & Analytics*. We note that this choice is arbitrary, and does not effect the reported results qualitatively. We also performed experiments on other variants of the oDesk datasets that considered different sets of task types.

Each worker in the oDesk has a profile with rich meta-data available via an API. In particular, the fields that are of particular interest to us include: *i*) the “skills” (a set of free-form text tags that workers can assign to themselves); *ii*) feedback score based on previous tasks completed; *iii*) number of hours worked; and *iv*) the top-level categories of the tasks completed by the workers, based on which the feedback score is aggregated. We crawled a sample of 200 workers by issuing a specific query². The skills in this query were chosen so as to ensure that the completed jobs by the workers in the retrieved list possibly have some overlap with the top-level categories of the M tasks, otherwise, this overlap would be low for a randomly retrieved list of workers. In realistic setting, this overlap is expected as workers bid for tasks based on their skills and job profile. The number of

¹<https://developers.odesk.com/>

²{‘hours’: ‘[100 TO 10000]’, ‘skills’: ‘cartooning OR machine-learning OR translation’}

hours worked was set to a minimum of 100 to ensure there is sufficient feedback available for the workers, given that the feedbacks are generally sparse.

We created the side-observation graphs as follows. We add an edge between two tasks o_x and o_z , *i.e.*, $\{o_x, o_z\} \in E_o$, if these two tasks belong to the same top-level category. In our setting, this would result in 3 disconnected cliques among the 10 tasks. For the workers, we computed the Jaccard's coefficient between the skills of any two workers. We add an edge between two workers w_x and w_z , *i.e.*, $\{w_x, w_z\} \in E_w$, if the Jaccard's coefficient between w_x and w_z is above a certain threshold (chosen to be 0.3 for the reported results). Next, we create the performance matrix from the feedback scores in a similar manner to the approach we took with the synthetic data. First, for each task o_j , we sampled Δ_j^{min} uniformly at random from range $[0.01, 0.5]$. Then, for a given worker w_i and task o_j , we look at the feedback score of w_i obtained in the historically completed tasks which belong to top-level category same as that of o_j . Note that this feedback score is in the range of $[0, 5]$ rating. When available, this feedback score is used for $\mu_{(i,j)}$, else feedback score is randomly sampled from $[0, 3]$. These feedbacks are then scaled to lie in the range $[\mu^{min}, \mu^{max} - \Delta_j^{min}]$, except for best worker for o_j , whose $\mu_{(i,j)}$ is set to μ^{max} . This process is repeated for each of the M tasks independently.

Results

We now discuss the findings from our experiments.

Varying budget and measuring precision. Figure 2(a) and Figure 3(a) shows the results for varying the average budget spent per worker/task pair, and how it leads to increased precision of the team selected by UEXPSELECT and AEXPSELECT. For these results, the average $\Delta_j^{min} = 0.25$, and is same for both the synthetic data (Figure 2(a)) and the oDesk data (Figure 3(a)). For both the datasets, AEXPSELECT shows significantly faster convergence towards selecting the optimal team. For instance, in Figure 2(a), AEXPSELECT achieved over 90% precision (getting the ϵ -optimal worker for 9 out of 10 tasks) at budget of $20 \cdot M \cdot N$, whereas UEXPSELECT requires substantially much more budget to achieve same precision. The difference in performance of AEXPSELECT or UEXPSELECT across synthetic and oDesk datasets is simply attributed to the different distribution of the workers' performances across the datasets. In particular, in the oDesk data, the performance values of the workers are more skewed towards higher values making it more challenging problem instance, in comparison to the synthetic data where the performance values are sampled uniformly.

Varying budget and measuring performance gap. Figure 2(b) shows an alternate view of the corresponding result in Figure 2(a). While Figure 2(a) reported 0/1 loss, Figure 2(b) reports the average of the actual performance gap of the best worker for a task in the output set compared to the best worker in full set.

Varying hardness of problem instance. In Figure 2(c), budget is kept fixed to 20 per worker/task pair, and average Δ_j^{min} , quantifying the hardness of the problem instance, is changed from 0.25 to 0.05. The gain of adaptive assignments

in AEXPSELECT compared to UEXPSELECT is consistent, though both the algorithms degrade in performance as expected.

Effect of exploiting side-observations. In Figure 3(a), there is no side-observation graph, ($E_o = \emptyset, E_w = \emptyset$), and the $\tilde{\gamma}_{G_{wo}}$ is simply equal to $M \cdot N$. In Figure 3(b), there is side-observation graph over the tasks as described in the data generation, however no graph is used over workers ($E_w = \emptyset$). The $\tilde{\gamma}_{G_{wo}}$ in this case as computed by greedy algorithm is equal to 600. Figure 3(c) shows results which considers side-observation graphs over both tasks and workers, with $\tilde{\gamma}_{G_{wo}} = 386$. Both the algorithms see a significant boost in terms of faster learning by exploiting the side-observations. Furthermore, we can see that the boost in performance by adding side-observations is more for UEXPSELECT compared to AEXPSELECT, as discussed during the theoretical performance analysis of the algorithms.

Conclusions and Future Work

We presented an algorithmic approach to tackle the challenge of the efficient hiring of teams of workers, as faced by recruiters for contract-based crowdsourcing. By casting these budgeted decision-theoretic problems as an instance of online learning for best action selection, we designed algorithms with PAC bounds, and further extended them to exploit the commonalities among the tasks and the workers. Our methodology and results present an interesting direction of continued research for the problem of hiring a team for contract-based crowdsourcing.

We see several interesting directions in which the current work can be extended. In particular, we used a simple notion of quantifying the optimality of the team. We see promise in extending the results to incorporate more complex relations among team members, such as the matching of task types within teams to balance the workload, capturing diminishing returns of growing teams, learning and representing costs associated with communication and coordination among people with different skills and abilities (including collaborative competency), and other combinatorial constraints, as an interesting direction for future work. Furthermore, we are interested in developing more realistic models of side-observations and performing real-world experiments using those models.

References

- [Bubeck, Munos, and Stoltz 2009] Bubeck, S.; Munos, R.; and Stoltz, G. 2009. Pure exploration in multi-armed bandits problems. In *Algorithmic Learning Theory (ALT)*, 23–37.
- [Buccapatnam, Eryilmaz, and Shroff 2014] Buccapatnam, S.; Eryilmaz, A.; and Shroff, N. B. 2014. Stochastic bandits with side observations on networks. In *SIGMETRICS*.
- [Caron et al. 2012] Caron, S.; Kveton, B.; Lelarge, M.; and Bhagat, S. 2012. Leveraging side observations in stochastic bandits. In *UAI*.
- [Cesa-Bianchi, Gentile, and Zappella 2013] Cesa-Bianchi, N.; Gentile, C.; and Zappella, G. 2013. A gang of bandits. In *NIPS*, 737–745.
- [Chen et al. 2014] Chen, S.; Lin, T.; King, I.; Lyu, M. R.; and Chen, W. 2014. Combinatorial pure exploration of multi-armed bandits. In *NIPS*, 379–387.
- [Difallah, Demartini, and Cudré-Mauroux 2013] Difallah, D. E.; Demartini, G.; and Cudré-Mauroux, P. 2013. Pick-a-crowd: tell me what you like, and i’ll tell you what to do. In *WWW*, 367–374.
- [Even-Dar, Mannor, and Mansour 2006] Even-Dar, E.; Mannor, S.; and Mansour, Y. 2006. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of Machine Learning Research* 7:1079–1105.
- [Fang and Tao 2014] Fang, M., and Tao, D. 2014. Networked bandits with disjoint linear payoffs. In *SIGKDD*, 1106–1115.
- [Gabillon et al. 2011] Gabillon, V.; Ghavamzadeh, M.; Lazaric, A.; and Bubeck, S. 2011. Multi-bandit best arm identification. In *NIPS*, 2222–2230.
- [Goel, Nikzad, and Singla 2014] Goel, G.; Nikzad, A.; and Singla, A. 2014. Mechanism design for crowdsourcing markets with heterogeneous tasks. In *HCOMP*.
- [Guha and Khuller 1998] Guha, S., and Khuller, S. 1998. Approximation algorithms for connected dominating sets. *Algorithmica* 20(4):374–387.
- [Ho and Vaughan 2012] Ho, C.-J., and Vaughan, J. W. 2012. Online task assignment in crowdsourcing markets. In *AAAI*.
- [Ho, Jabbari, and Vaughan 2013] Ho, C.-J.; Jabbari, S.; and Vaughan, J. W. 2013. Adaptive task assignment for crowd-sourced classification. In *ICML*, 534–542.
- [Ho, Slivkins, and Vaughan 2014] Ho, C.-J.; Slivkins, A.; and Vaughan, J. W. 2014. Adaptive contract design for crowdsourcing markets: bandit algorithms for repeated principal-agent problems. In *EC*, 359–376. ACM.
- [Hoeffding 1963] Hoeffding, W. 1963. Probability inequalities for sums of bounded random variables. *Journal of the American statistical association* 58(301):13–30.
- [Kalyanakrishnan et al. 2012] Kalyanakrishnan, S.; Tewari, A.; Auer, P.; and Stone, P. 2012. PAC subset selection in stochastic multi-armed bandits. In *ICML*, 655–662.
- [Lai and Robbins 1985] Lai, T. L., and Robbins, H. 1985. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics* 6(1):4–22.
- [Mannor and Shamir 2011] Mannor, S., and Shamir, O. 2011. From bandits to experts: On the value of side-observations. In *NIPS*, 684–692.
- [Shahaf and Horvitz 2010] Shahaf, D., and Horvitz, E. 2010. Generalized task markets for human and machine computation. In *AAAI*.
- [Singla and Krause 2013] Singla, A., and Krause, A. 2013. Truthful incentives in crowdsourcing tasks using regret minimization mechanisms. In *WWW*, 1167–1178.
- [Valiant 1984] Valiant, L. G. 1984. A theory of the learnable. *Communications of the ACM* 27(11):1134–1142.
- [Wang, Viswanathan, and Bubeck 2013] Wang, T.; Viswanathan, N.; and Bubeck, S. 2013. Multiple identifications in multi-armed bandits. In *ICML*, 258–265.
- [Zhang et al. 2011] Zhang, H.; Horvitz, E.; Miller, R. C.; and Parkes, D. C. 2011. Crowdsourcing general computation. In *HCOMP Workshop, CHI*.
- [Zhou, Chen, and Li 2014] Zhou, Y.; Chen, X.; and Li, J. 2014. Optimal pac multiple arm identification with applications to crowdsourcing. In *ICML*.